

A category theoretic approach to unital quantum channel convexity and Birkhoff's theorem

Ian T. Durham*

Department of Physics, Saint Anselm College, Manchester, NH 03102

(Dated: December 30, 2009)

Birkhoff's Theorem states that doubly stochastic matrices are convex combinations of permutation matrices. Quantum mechanically these matrices are doubly stochastic channels, i.e. they are completely positive maps preserving both the trace and the identity. We expect these channels to be convex combinations of unitary channels and yet it is known that some channels cannot be written that way. Recent work has suggested that n copies of a single channel might approximate a mixture (convex combination) of unitaries. In this paper we show that $n(n + 1)/2$ copies of a symmetric unital quantum channel may be arbitrarily-well approximated by a mixture (convex combination) of unitarily implemented channels. In addition, we prove that any extremal properties of a channel are preserved over n (and thus $n(n + 1)/2$) copies. The result has the potential to be completely generalized to include non-symmetric channels.

I. INTRODUCTION AND BACKGROUND

There is a famous theorem attributed to Garrett Birkhoff that states that doubly stochastic matrices are convex combinations of permutation matrices. In the quantum context, doubly stochastic matrices become doubly stochastic channels, i.e. completely positive maps preserving both the trace and the identity. Quantum mechanically we understand the permutations to be the unitarily implemented channels. That is, we expect doubly stochastic quantum channels to be convex combinations of unitary channels. Unfortunately it is well-known that some quantum channels *cannot* be written that way [1, 2]. Recent work has suggested that n copies of a single channel might approximate a mixture (convex combi-

*idurham@anselm.edu

tion) of unitaries [3, 4, 5, 6]. In this article we prove a slightly stronger result for symmetric unital quantum channels via the following theorem.

Theorem I.1 *Given a symmetric unital quantum channel that maps from ρ to $T(\rho)$, we may approximate $n(n+1)/2$ copies of such a channel arbitrarily well with a mixture (convex combination) of unitarily implemented channels.*

In addition, we demonstrate that extremal properties of unital quantum channels are maintained over $n(n+1)/2$ copies and that such channels behave classically in the limit of classical information.

We prove the aforementioned theorem in several stages. First we discuss convexity beginning with a review of the classical version of Birkhoff's theorem. We then discuss convexity as it relates to unital quantum channels and prove that n (and thus $n(n+1)/2$) copies of a unital quantum channel T retain any extremal properties T might possess, where n has no restrictions. After a brief review of basic category theory, we demonstrate that a single symmetric unital quantum channel along with the associated Hilbert space is a category with one object. By definition this also means that it is a group. Cayley's theorem then implies that it is also isomorphic to some permutation group. We then construct a larger category consisting of $n(n+1)/2$ copies of our unital quantum channel with this category (which consists of groups) being isomorphic to a set of permutation groups. Finally we apply a group theoretic result to show that these permutation groups are, in fact, unitary.

II. CONVEXITY

We begin by reviewing Birkhoff's theorem (also known as the Birkhoff-von Neumann theorem). It may be stated in a number of ways and our presentation follows closely that given in Steele[7]. Given a permutation $\sigma \in S_n$, the permutation matrix that is associated with σ is the $n \times n$ matrix $P_\sigma(j, k)$, where $1 \leq j, k \leq n$, whose entries are given by

$$P_\sigma(j, k) = \begin{cases} 1 & \text{if } \sigma(j) = k \\ 0 & \text{otherwise.} \end{cases}$$

An $n \times n$ doubly stochastic matrix D is a square matrix whose elements are real and whose rows and columns sum to unity. Given such a matrix, there exist nonnegative weights

$\{w_\sigma : \sigma \in S_n\}$ such that

$$\sum_{\sigma \in S_n} w_\sigma = 1 \quad \text{and} \quad \sum_{\sigma \in S_n} w_\sigma P_\sigma = D. \quad (1)$$

In other words, every doubly stochastic matrix is an average (convex combination of) permutation matrices. The set of such matrices of order n is said to form the convex hull of permutation matrices of the same order where the latter are the vertices (extreme points) of the former. It superficially appears as if this ought to translate quite easily into the quantum world, particularly if we were to assume that the nonnegative weights were actually probabilities. Unfortunately that is not the case.

In the quantum world we are particularly interested in what a matrix like D can *do*. In other words, if we treat it like a map or a ‘black box’[23] we wish to note what sort of an output we get given a specific input. Generally we are interested in completely positive (CP) maps which are maps preserving both trace and identity (i.e. and identity preserving CPTP map). They are represented by square matrices of dimension d and are known as unital quantum channels. They can be written as a *Kraus decomposition* as

$$T(\rho) = \sum_i A_i \rho A_i^\dagger \quad \text{where} \quad \sum_i A_i^\dagger A_i = \mathbb{1} \quad (2)$$

and $T(\mathbb{1}) = \mathbb{1}$. The set of all quantum channels on \mathcal{M}_d is convex and compact meaning it may be decomposed as

$$T = \sum_i p_i T_i \quad (3)$$

where the p ’s are probabilities and the T_i ’s are *extremal* unital channels, that is channels that may not be further decomposed [5]. Channels with a single Kraus operator are pure channels and the extremal points in the convex set of channels are precisely the pure channels. Here T represents the set of *all* channels on the particular space, not necessarily copies of the same one, i.e. the T_i may not represent the same channel. T also has additional restrictions not possessed by D .

A. Preservation of extremal properties for unital quantum channels

The preservation of unitality over tensor products is a basic property of the tensor product of algebras. As such this implies that the trace-preserving property is itself preserved, i.e.

$$\sum_{i_1, \dots, i_n} (A_{i_1} \otimes \dots \otimes A_{i_n})^\dagger (A_{i_1} \otimes \dots \otimes A_{i_n}) = \text{id} \quad (4)$$

However, it is not clear that the extremal properties, if possessed, are also preserved. Specifically, Landau and Streater have proven a theorem that if T is unital, it is extremal if and only if the set

$$\left\{ A_k^\dagger A_l \oplus A_l A_k^\dagger \right\}_{k,l \dots N} \quad (5)$$

is linearly independent [1]. Thus, as we take a tensor product of T with itself, if we wish for the resultant channel to remain extremal we require that the set

$$\{(A_{k_1}^\dagger \otimes A_{k_2}^\dagger)(A_{l_1} \otimes A_{l_2}) \oplus (A_{l_1} \otimes A_{l_2})(A_{k_1}^\dagger \otimes A_{k_2}^\dagger)\}_{k_1, k_2, l_1, l_2} \quad (6)$$

remain linearly independent. Generalized over n tensor products we require that

$$\left\{ \left(A_{k_1}^\dagger \otimes \cdots \otimes A_{k_n}^\dagger \right) (A_{l_1} \otimes \cdots \otimes A_{l_n}) \oplus (A_{l_1} \otimes \cdots \otimes A_{l_n}) \left(A_{k_1}^\dagger \otimes \cdots \otimes A_{k_n}^\dagger \right) \right\}_{k, l \dots N} \quad (7)$$

be linearly independent. This is amounts to showing that equation (7) is of the same basic form as equation (5). In other words, we require that

$$\begin{aligned} \left(A_{k_1}^\dagger \otimes \cdots \otimes A_{k_n}^\dagger \right) (A_{l_1} \otimes \cdots \otimes A_{l_n}) &= \left(A_j^\dagger A_m \right)_{j, m \dots N} \\ (A_{l_1} \otimes \cdots \otimes A_{l_n}) \left(A_{k_1}^\dagger \otimes \cdots \otimes A_{k_n}^\dagger \right) &= \left(A_m A_j^\dagger \right)_{j, m \dots N}. \end{aligned} \quad (8)$$

In the case that the Kraus operators are symmetric this can be easily accomplished via a Cholesky decomposition in which any symmetric, square, positive definite matrix can be decomposed into the product of a lower triangular matrix and its conjugate transpose (there is also a version that uses an upper triangular matrix) while preserving linear independence [8]. In other words, as long as the matrices on the left side of equation (8) are symmetric, Cholesky decomposition may be used to find equations of the form required by the right side of equation (8). Even though Kraus operators are not necessarily unique for a given quantum operation, all systems of Kraus operators that represent the same quantum operation are related via a unitary transformation (see Theorem 8.2 in [9]).

Suppose we have a set of Kraus operators for a given unital quantum channel that is known to be extremal. Since extremality is a unique metric property there will be a unique value for the fidelity of the channel if it is extremal. Since fidelity is invariant under unitary transformations and the fidelity associated with an extremal channel is unique, extremality is preserved under such transformations. Thus, even though the Kraus operators may not be unique for a given extremal channel, the Cholesky decomposition as applied in equation (8) will preserve extremality for the set of all extremal Kraus operators for that channel.

Note that while Cholesky decomposition only works for symmetric matrices, this doesn't severely limit the relevance of these results since there do exist symmetric channels that violate Birkhoff's theorem. An example of one such channel is the 'A not Q' channel given in [2]. Since this channel is diagonal it is symmetric and thus Cholesky decomposition may be used to prove (8) and, by the above argument regarding fidelity, extremality is preserved.

In general, we expect that most quantum channels are symmetric if they are to preserve identity, though this is not a foregone conclusion (i.e. there are matrices that could satisfy equation (2), preserving identity, that are not symmetric, though we doubt any such matrices are physically significant). At any rate, it is presently unknown (at least to this author) whether a method exists to similarly decompose *non-symmetric* (or non-diagonalizable) Kraus operators for a quantum channel while retaining the extremal properties in the asymptotic limit, i.e. replicating equation (8).

In any case, we have at least shown that *symmetric* channels, including certain troublesome channels (e.g. the 'A not Q' channel mentioned above), retain their extremal characteristics in the asymptotic limit. This is an important step in verifying the suggested quantum extension of Birkhoff's theorem.

III. A REVIEW OF CATEGORIES

We now pause for a brief review of category theory. (The reader is encouraged to consult a standard text on the subject [10, 11] as well as materials introducing it to quantum theory [12, 13, 14, 15].) The following definitions closely follow those given in Awodey[11].

Category A *category* is a mathematical structure that consists of *Objects*: A, B, C, \dots and *Arrows*: f, g, h, \dots with the properties,

- for each arrow f there are given objects $\text{dom}(f)$ and $\text{cod}(f)$ called the *domain* and *codomain* of f respectively. We write $f : A \rightarrow B$ to indicate that $A = \text{dom}(f)$ and $B = \text{cod}(f)$;
- given arrows $f : A \rightarrow B$ and $g : B \rightarrow C$ where $\text{cod}(f) = \text{dom}(g)$ there is given an arrow $g \circ f : A \rightarrow C$ called the *composite* of f and g ;
- for each object A there is given an arrow $1_A : A \rightarrow A$ called the *identity arrow* of A .

These must satisfy

- **Associativity** $h \circ (g \circ f) = (h \circ g) \circ f$ for all $f : A \rightarrow B, g : B \rightarrow C, h : C \rightarrow D$;
- **Unit** $f \circ \mathbb{1}_A = f = \mathbb{1}_B \circ f$ for all $f : A \rightarrow B$.

Functor A *functor* $F : \mathbf{C} \rightarrow \mathbf{D}$ between categories \mathbf{C} and \mathbf{D} is a mapping from objects to objects and arrows to arrows such that

- $F(f : A \rightarrow B) = F(f) : F(A) \rightarrow F(B)$,
- $F(g \circ f) = F(g) \circ F(f)$,
- $F(\mathbb{1}_A) = \mathbb{1}_{F(A)}$.

Definition In any category \mathbf{C} , an arrow $f : A \rightarrow B$ is called an isomorphism if there is an arrow $g : B \rightarrow A$ in \mathbf{C} such that $g \circ f = \mathbb{1}_A$ and $f \circ g = \mathbb{1}_B$.

In category theory, tensor products form a bifunctor from the category of vector spaces to itself and are covariant in each argument. As such they may be decomposed as

$$T(\rho) \otimes T(\rho) = (T \otimes T)(\rho \otimes \rho). \quad (9)$$

That is, if T is a linear map $T : R \rightarrow R$ with $\rho \in R$, then $T \otimes T$ is a linear map $T \otimes T : R \otimes R \rightarrow R \otimes R$ and $\rho \otimes \rho \in R \otimes R$. This is easily generalized to n copies. In other words, $\rho^{\otimes n}$ is the input state to n channels. The definition of functors requires that they preserve identity morphisms [10, 16]. It is clear that a unital channel is an identity morphism since $T(\mathbb{1}) = \mathbb{1}$ and thus n copies of such a channel should preserve this property. In other words, $T^{\otimes n}(\mathbb{1}^{\otimes n}) = \mathbb{1}^{\otimes n}$.

IV. CATEGORY AND GROUP STRUCTURE OF QUANTUM CHANNELS

A unital quantum channel is a mapping between Hilbert spaces, $\Phi : L(\mathcal{H}_A) \rightarrow L(\mathcal{H}_B)$, where $L(\mathcal{H}_i)$ is the family of operators on \mathcal{H}_i . The operator spaces can be interpreted as C^* -algebras and thus we can also view the channel as a mapping between C^* -algebras, $\Phi : \mathcal{A} \rightarrow \mathcal{B}$. Quantum channels can carry classical information as well. An example of such a channel would be $\Phi : L(\mathcal{H}_A) \otimes C(X) \rightarrow L(\mathcal{H}_B)$ where $C(X)$ is the space of continuous

functions on some set X and is also a C^* -algebra. In other words, whether or not classical information is processed by the channel, it (the channel) is a mapping between C^* -algebras. Note, however, that these are not necessarily the same C^* -algebras. Since the channels are represented by square matrices, the input and output C^* -algebras must have the same dimension, d . Thus we can consider them both subsets of some d -dimensional C^* -algebra, \mathcal{C} , i.e. $\mathcal{A} \subset \mathcal{C}$ and $\mathcal{B} \subset \mathcal{C}$. Thus a unital quantum channel is a mapping from \mathcal{C} to itself (the need for a single object will become apparent in a moment). As such we propose the following:

Proposition IV.1 *A unital quantum channel given by $t : L(\mathcal{H}_\rho) \rightarrow L(\mathcal{H}_{T(\rho)})$, together with the d -dimensional C^* -algebra, \mathcal{C} , on which it acts, forms a category we call $\mathbf{Chan}(d)$.*

Proof Our object in this case is \mathcal{C} (the C^* -algebra) while our arrow is t (the channel). The existence of a domain and codomain of t are trivially obvious. We must show compositeness. Consider the unital quantum channels

$$\begin{aligned} r : L(\mathcal{H}_\rho) &\rightarrow L(\mathcal{H}_\sigma) & \text{where} & \sigma = \sum_i A_i \rho A_i^\dagger \\ t : L(\mathcal{H}_\sigma) &\rightarrow L(\mathcal{H}_\tau) & \text{where} & \tau = \sum_j B_j \sigma B_j^\dagger \end{aligned}$$

where the usual properties of such channels are assumed (e.g. trace preserving, etc.). We form the composite $t \circ r : L(\mathcal{H}_\rho) \rightarrow L(\mathcal{H}_\tau)$ where

$$\begin{aligned} \tau &= \sum_j B_j \left(\sum_i A_i \rho A_i^\dagger \right) B_j^\dagger \\ &= \sum_{i,j} B_j A_i \rho A_i^\dagger B_j^\dagger \\ &= \sum_k C_k \rho C_k^\dagger \end{aligned} \tag{10}$$

. Since A and B are summed over separate indices the trace-preserving property is maintained, i.e. $\sum_k C_k^\dagger C_k = \mathbb{1}$ For a similar methodology, see [17].

There clearly exists an identity arrow in at least one instance due to the nature of unital quantum channels, i.e. by definition $T(\mathbb{1}) = \mathbb{1}$. We may take the completely *general* identity arrow

$$\mathbb{1}_\rho : L(\mathcal{H}_\rho) \rightarrow L(\mathcal{H}_\rho)$$

to be the time evolution of the state ρ in the *absence* of any unital quantum channel. Since this definition is suitably general we have that

$$t \circ \mathbb{1}_A = t = \mathbb{1}_B \circ t \quad \forall t : A \rightarrow B$$

and we are thus left to prove the associativity of composition.

Consider the three unital quantum channels $r : L(\mathcal{H}_\rho) \rightarrow L(\mathcal{H}_\sigma)$, $t : L(\mathcal{H}_\sigma) \rightarrow L(\mathcal{H}_\tau)$, and $v : L(\mathcal{H}_\tau) \rightarrow L(\mathcal{H}_v)$ where $\sigma = \sum_i A_i \rho A_i^\dagger$, $\tau = \sum_j B_j \sigma B_j^\dagger$, and $v = \sum_k C_k \tau C_k^\dagger$. We have

$$\begin{aligned} v \circ (t \circ r) &= v \circ \left(\sum_{i,j} B_j A_i \rho A_i^\dagger B_j^\dagger \right) = \sum_k C_k \left(\sum_{i,j} B_j A_i \rho A_i^\dagger B_j^\dagger \right) C_k^\dagger \\ &= \sum_{i,j,k} C_k B_j A_i \rho A_i^\dagger B_j^\dagger C_k^\dagger = \sum_{i,j,k} C_k B_j \left(A_i \rho A_i^\dagger \right) B_j^\dagger C_k^\dagger \\ &= \left(\sum_{i,j,k} C_k B_j \tau B_j^\dagger C_k^\dagger \right) \circ r = (v \circ t) \circ r \end{aligned}$$

and thus we have associativity. \square

We have intentionally defined our category such that we have a single object, \mathcal{C} , rather than two objects, ρ and $T(\rho)$. A category with one object is known as a *monoid*. Restricting ourselves to monoids allows us to make use of the following definition.

Definition A group G is a category with one object, in which every arrow is an isomorphism [11]. Conversely, any category with a single object whose arrows are invertible is a group under composition [18].

Note that the invertibility of the arrows in category theoretic terms does not necessarily mean that the channel itself is invertible, i.e. unitary. If it were, these steps would be unnecessary since there are much simpler ways to show that unitary channels stay unitary in the asymptotic limit. The invertibility here, however, is at the category-theoretic level as is the isomorphism. In fact, as Awodey points out, in many cases the *only* definition that makes sense is the abstract, category-theoretic one. Thus, while the unital quantum channel itself may not be invertible, a category formed with it might be. Physically, an example of this would be the following. Suppose we have a non-polarized light beam that passes through a polarizer. This could be modeled as our quantum channel and is clearly not invertible, i.e. we can't pass the polarized light back through the polarizer to depolarize

it. However, we *could* pass the polarized light through an optical depolarizer. In a category theoretic sense we would still only have one object, though we would now have two arrows that satisfy the above definition of a category-theoretic isomorphism. As such, our category **Chan**(d), is a group. This allows us to make use of a powerful theorem of group theory known as Cayley's theorem.

Theorem IV.2 (Cayley) *Every group is isomorphic to a permutation group.*

There are numerous presentations and proofs of this theorem. Since it is well-known, we refer the reader to one of the following for a proof [11, 19, 20, 21].

This means that a unital quantum channel and the associated C^* -algebra *is isomorphic to a permutation group*. We might think we can simply say now that n copies of this channel would be isomorphic to a set of permutation groups and leave it at that. The problem is that the dimension of the C^* -algebra is dependent on the value of n because two copies of a channel are connected by a tensor product.

Let us then define **Chan**(d) to be a category with a d -dimensional C^* -algebra as the sole object. So, for example, suppose our basic channel is associated with a two-dimensional C^* -algebra. Its category would be **Chan**(2). Two copies of this channel exist are associated with a four-dimensional C^* -algebra and so their category would be **Chan**(4). We can generalize this a bit more by defining **Chan**(d^n) to be a category with a d^n -dimensional C^* -algebra.

Categories are like groups in that one can have a category made up of categories. Thus we define the category **Chan** with objects **Chan**(d), **Chan**(d^2), ..., **Chan**(d^n) and the arrow \otimes . Since tensor products are known to be associative, we can move between the objects in our new category fairly easily. For example

$$[\mathbf{Chan}(d) \xrightarrow{\otimes} \mathbf{Chan}(d^2)] \xrightarrow{\otimes} \mathbf{Chan}(d^3) = \mathbf{Chan}(d) \xrightarrow{\otimes} [\mathbf{Chan}(d^2) \xrightarrow{\otimes} \mathbf{Chan}(d^3)]. \quad (11)$$

Note in this case that the arrow is not an isomorphism. This is fine since the isomorphism criterion was only needed in order to utilize Cayley's theorem to ensure that each **Chan**(d^n) was isomorphic to a permutation group. Also notice that the category **Chan** contains $n(n+1)/2$ copies of our original channel rather than simply n . We will discuss the importance of this after proving the main theorem of this article.

We are now in a position to prove the central theorem of our paper.

Theorem IV.3 *Given a symmetric unital quantum channel that maps from ρ to $T(\rho)$, we may approximate $n(n+1)/2$ copies of such a channel arbitrarily well with a mixture (convex combination) of unitarily implemented channels.*

Proof By Cayley's theorem we've shown that each category, $\mathbf{Chan}(d^n)$, which is also a group, is individually isomorphic to some permutation group. As a group, $\mathbf{Chan}(d^n)$ has infinite cardinality since we're invoking the C^* -algebra for the entire space of inputs and outputs. Thus, since Cayley's theorem assumes an exact isomorphism, the permutation group must also have infinite cardinality. However, each of these infinite permutation groups is ε -isomorphic to some *finite* permutation group. Note that any representation of a permutation group that is finite and compact may be considered unitary where a representation is a set of operators representing the permutations, i.e. the operators that produce the permutations [22]. Thus *representations* of finite and compact permutation groups are necessarily unitary. So, individually, the channels are ε -isomorphic to something that has a unitary representation. As $n \rightarrow \infty$, the cardinality of the *finite* permutation groups increases and thus the isomorphism approaches exact. The approach is asymptotic as long as n remains finite which it must in order for the representation to be unitary. Since the category \mathbf{Chan} , while containing n copies of $\mathbf{Chan}(d^n)$, actually contains $n(n+1)/2$ copies our unital quantum channel, it is $n(n+1)/2$ copies of that channel that approaches a representation of a set of unitarily represented groups approximately well. Regarding the term *arbitrarily well*, we interpret this to mean that we may choose *any* value of ε , no matter how small, and still find a better approximation. This fact remains fundamentally true since n can increase indefinitely while remaining finite. \square

We note that there are different levels of isomorphism present throughout this proof. Some are on a category theoretic-level while others are at a group-theoretic level. This is an important distinction that can affect the interpretation of these results. As an analogy, a Euro and a Dollar are related by dint of the fact that both are the basis of a particular currency. Conversely, within, say, the American system, one could possess either a paper dollar or a dollar coin. This is an entirely different level of relation.

V. DISCUSSION

These results represent the most general version of Birkhoff’s theorem in the asymptotic limit yet discovered, though we readily admit that the level of abstraction used (including the representation theoretic aspects in the proof of the final theorem) may leave doubts in some readers’ minds. On the other hand, when considering the physical extension of these abstractions, we believe the result captures the essence of the suggested asymptotic limit to Birkhoff’s theorem. In particular, considering the fact that we have an isomorphism to a *representation* in the final proof, we note that, physically, Birkhoff’s theorem only implies that we may, in essence, *represent* n copies of a unital quantum channel with a convex combination of unitarily implemented ones. In other words, we could build something in a lab that approximates n copies of some, perhaps complicated, unital quantum channel. So we interpret ‘to approximate’ as meaning essentially the same thing as ‘to represent.’ This is, of course, simply an interpretation and we believe that it would greatly benefit from experimental results. In addition, as we have noted, there are different levels of isomorphism being used here and it is important to note the differences when interpreting the results.

The results are of particular importance for physical realizations of certain ‘problematic’ unital quantum channels, i.e. those that can’t be approximated by unitary channels individually. It implies that we can approximate n copies of them unitarily allowing us to study their behavior and approximate that behavior in realistic quantum information processing systems. In particular, a quantum channel allows for perfect environment-assisted error correction if and only if it is a mixture of unitaries. As such, this result implies we can approximate a system with perfect environment-assisted error correction with a set of unitarily implemented channels for which many relatively simple physical manifestations already exist.

In addition we note that this result is stronger than the original suggestion that the best we could approximate was n copies of a given channel. Our result implies we may actually approximate $n(n + 1)/2$ copies. While technically larger than n in the asymptotic limit, the number of computational steps required to reach a very large number of copies of our channel is *less*. This is because each individual category is represented by a single unitarily implemented channel and thus n such channels approximate $n(n + 1)/2$ unital quantum channels.

Finally, we note that the restriction to symmetric channels is in order to preserve any extremal characteristics. If a method is found that could decompose a non-symmetric matrix into a form similar to the right-hand side of equation (8), then our result would be completely general.

Acknowledgments

We thank Aram Harrow and Steve Shea for in depth and on-going discussions of this problem. Additionally, we thank Michael Ben-Or and three anonymous reviewers who provided very useful comments. We also acknowledge several informal conversations concerning various aspects of this paper with Dave Bacon, Winton Brown, Lorenza Viola, and attendees of one of the physics/mathematics seminars at Saint Anselm College.

[1] L. Landau and R. Streeter, *Journal of Linear Algebra and Its Applications* **193**, 107 (1993).

[2] M. Gregoratti and R. Werner, *Journal of Modern Optics* **50**, 915 (2003).

[3] A. Winter, *On environment-assisted capacities of quantum channels*, <http://arxiv.org/abs/quant-ph/0507045> (2005).

[4] J. A. Smolin, F. Verstraete, and A. Winter, *Physical Review A* **72**, 052317 (2005).

[5] C. B. Mendl and M. M. Wolf, *Unital quantum channels - convex structure and revivals of birkhoff's theorem*, <http://arxiv.org/abs/0806.2820> (2008).

[6] J. Watrous, *Mixing doubly stochastic quantum channels with the completely depolarizing channel*, <http://arxiv.org/abs/0807.2668> (2008).

[7] J. M. Steele, *The Cauchy-Schwarz Master Class* (Cambridge University Press, Cambridge, 2004).

[8] D. S. Watkins, *Fundamentals of Matrix Decomposition* (John Wiley and Sons, New York, 2002).

[9] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).

[10] S. MacLane, *Categories for the working mathematician* (Springer, Berlin, 1971).

[11] S. Awodey, *Category Theory*, no. 49 in Oxford Logic Guides (Clarendon Press, Oxford, 2006).

- [12] B. Coecke, *Advanced Studies in Mathematical Logic* **30**, 45 (2006).
- [13] S. Abramsky and B. Coecke, in *Proceedings of the 19th IEEE conference on Logic in Computer Science (LiCS'04)* (IEEE Computer Society Press, Los Alamitos, 2004).
- [14] S. Abramsky and B. Coecke, in *Handbook of Quantum Logic and Quantum Structures, Volume II*, edited by K. Engesser, D. Gabbay, and D. Lehmann (Elsevier, Amsterdam, 2007).
- [15] J. Baez, in *Structural Foundations of Quantum Gravity*, edited by S. French, D. Rickles, and J. Saatsi (Oxford University Press, Oxford, 2006), pp. 240–265.
- [16] F. Borceux, *Handbook of Categorical Algebra: Volume 1, Basic Category Theory*, Encyclopedia of Mathematics and its Applications (Cambridge University Press, Cambridge, 1994).
- [17] A. Nayak and P. Sen, *Invertible quantum operations and perfect encryption of quantum states*, <http://arxiv.org/abs/0605041> (2006).
- [18] S. MacLane and G. Birkhoff, *Algebra, 3rd edition* (American Mathematical Society, Providence, 1999).
- [19] D. Wallace, *Groups, Rings, and Fields* (Springer, Berlin, 1998).
- [20] J. Beachy and W. Blair, *Abstract Algebra, Third Edition* (Waveland Press, Long Grove, 2006).
- [21] M. Armstrong, *Groups and Symmetry* (Springer, Berlin, 1997).
- [22] A. Böhm, *Quantum Mechanics: Foundations and Applications* (Springer, Berlin, 1993), 3rd ed.
- [23] Pedagogically we often make use of an old-fashioned meat grinder as a visual cue, inspired in part by an old *Monty Python* animated skit. We call it our ‘quantum meat grinder.’